


CONNEXIONS



The Interoperability Report

June 1991 Volume 5, No. 6

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

MIB Development.....	2
SNMP Security.....	12
The NREN Takes Shape.....	17
UK IP Activities.....	20
Psycoloquy.....	24
Letter to the Editor.....	26
NNSC Information.....	27
Book Review.....	28
Announcements.....	29

From the Editor

Network management continues to be an important topic in inter-networking. Several proprietary management systems exist, and OSI continues to develop standards in this arena, but if you want multi-vendor solutions today the *Simple Network Management Protocol* (SNMP) is the obvious answer. This month we bring you two articles related to SNMP. First, Bob Stewart discusses the development and integration of a *Management Information Base* (MIB), a critical component in SNMP. Second, Keith McCloghrie, Chuck Davin and Jim Galvin describe recently defined enhancements to SNMP to provide security and authentication.

Much has already been said and written about plans for a new *National Research and Education Network* (NREN). We asked Mike Roberts from EDUCOM to give us an overview of the NREN, and provide us with a "legislative update."

Across the Atlantic, the UK academic community has been looking at ways to provide support for the Internet suite of protocols ("TCP/IP") on the JANET network. JANET has been running its own protocol suite—known as the *Coloured Books*—for many years, and plans for an eventual transition to OSI. Therefore, the use of the Internet protocols turns out to be somewhat of a political "hot potato." Jon Crowcroft of UCL explains how TCP/IP is being offered (with certain constraints) in the *Shoestring* pilot project.

Computer networks are providing new ways for researchers to communicate. Bulletin boards and electronic mailing lists have been commonplace for decades, but a relatively new development is the concept of *electronic publishing* of refereed journals. One such journal, *PSYCOLOQUY*, has received critical acclaim. Stevan Harnad of Princeton University describes this project on page 24.

The *NSF Network Service Center* (NNSC) has produced several information "products" for use in the Internet community. Two of these, *The Internet Resource Guide* and a *HyperCard Tour of the Internet* are described by Karen Roubicek on page 27. Additionally, the NNSC publishes the *The Internet Manager's Phone Book*. The data in this directory was compiled from a wide variety of sources at the request of members of the IETF. It had been observed that it was often difficult for network managers to locate the proper people to call when debugging network problems. Contact information is scattered over several databases and these databases are often not accessible during network outages. The phone book attempts to address these issues by consolidating information in various databases, and providing a hardcopy which cannot be affected by network problems. For more information about the phone book, contact the NNSC: npsc@npsc.nsf.net • 617-873-3400.

Development and Integration of a Management Information Base

by Bob Stewart, Xyplex

Abstract With the Internet's SNMP as the primary example, and references to OSI's CMIP, this article positions the MIB concept in Network Management. It examines the growth of the SNMP MIB from first version through the current explosion of standard experimental and private MIBs. It discusses the reason, means, and problems of that growth, and propose a solution to the some of the problems.

Introduction A protocol to carry monitoring and control information is necessary to realize distributed network management, but the process of defining what the protocol carries encompasses far more effort and difficulty. Following, we will examine a process by which such information, termed a *Management Information Base* (MIB), is defined and provided to the network manager, using the *Simple Network Management Protocol* (SNMP) as the example. The discussion comprises the following major sections:

- Design—the origin, philosophy, and structure of SNMP.
- Development—the principles and process of MIB development.
- Future—the projected evolution of network management protocols and a proposal for improved interoperability.

History In 1988, recognizing the need for interoperable, distributed network management, the *Internet Activities Board* (IAB), overseeing body for the Internet protocol suite (commonly known as TCP/IP), commissioned work on the necessary protocol definitions [1]. The IAB chose a short-term plan for a simple protocol, with possible longer term replacement by an international standard protocol. To facilitate that evolution, the protocols were to share basic definitions.

The short term part of the plan lead to creation of SNMP, based heavily on the already-proven *Simple Gateway Monitoring Protocol* (SGMP) [3]. The long term part of the plan was to use the *Common Management Information Protocol* (CMIP) as defined by the *International Organization for Standardization* (ISO) in the *Open Systems Interconnection* (OSI) framework. For early Internet implementations, CMIP was to use the *Transmission Control Protocol* (TCP) as its transport mechanism, a combination known as "CMIP over TCP/IP" (CMOT) [4].

Work on SNMP progressed rapidly, but the CMOT work did not. Furthermore, attempts to coordinate the two paths were impeding progress, so in 1989 the IAB decoupled the efforts [2]. At this point, SNMP assumed the position of long-term solution. CMOT work continues in the Internet community, and CMIP work proceeds in the international community, but by far the most complete, interoperable network management implementation is found in SNMP. At this time, SNMP's base documents [5, 6, 7] compose a standard, recommended Internet protocol.

Philosophy The SNMP design philosophy is of major importance, as it drastically affects the protocol structure and evolution. The philosophy's two overriding principles are *simplicity* and *extensibility* [5]. Not only can we discuss the intent of these goals, but experience since they were chosen allows us to consider their merit and realization.

Simplicity is important to encourage implementation and foster correctness. Distributed network management is not possible unless the systems to be managed implement the protocol. Such implementations will not be present with the necessary ubiquity if such an implementation is too costly. Furthermore, they will not be useful if they are not robust and accurate. SNMP is therefore strongly biased toward simplicity in the managed systems, which far outnumber the managing systems.

For the most part, the design met the goal of simplicity, although simplicity in one area can lead to complexity in others (discussed later with regard to SNMP as a protocol). Such specifics notwithstanding, the widespread implementation of SNMP, versus the relative lack of practical CMIP implementations, indicates that SNMP's simplicity helped encourage its acceptance, while CMIP may have been hindered by its more extensive capability, but resulting complexity.

Extensibility is necessary to allow immediate, useful implementation while assuming that all necessary capabilities cannot be included at the outset. SNMP succeeds reasonably well in extensibility of MIB design, as shown by the proliferation of MIB extensions. It has been less effective in incorporating certain functions, such as actions to be taken, or efficient manipulation of large, complex databases. Nevertheless, this does not necessarily imply a shortcoming on the part of SNMP's designers, as the additional complexity accompanying such capabilities would most likely have violated the requirement for simplicity, resulting in a far more disastrous failure to be accepted as a standard. Furthermore, the "MIB explosion" was inevitable, as availability of interoperable, distributed network management lead to raised consciousness and expectations which turned into demands on commercial system vendors for extensive SNMP management capabilities.

Structure

Documentation of the SNMP design comprises three major components:

- The generic Structure of Management Information (SMI) [5].
- The Management Information Base [6].
- The protocol itself [7].

Overall, the approach to distributed network management follows the now-classic client-server model as shown in Figure 1. The client runs at the *managing system*. It makes requests and is typically called the *Network Management System* (NMS) or *Network Operation Center* (NOC). The server is in the *managed system*. It executes requests and is called the *agent*.

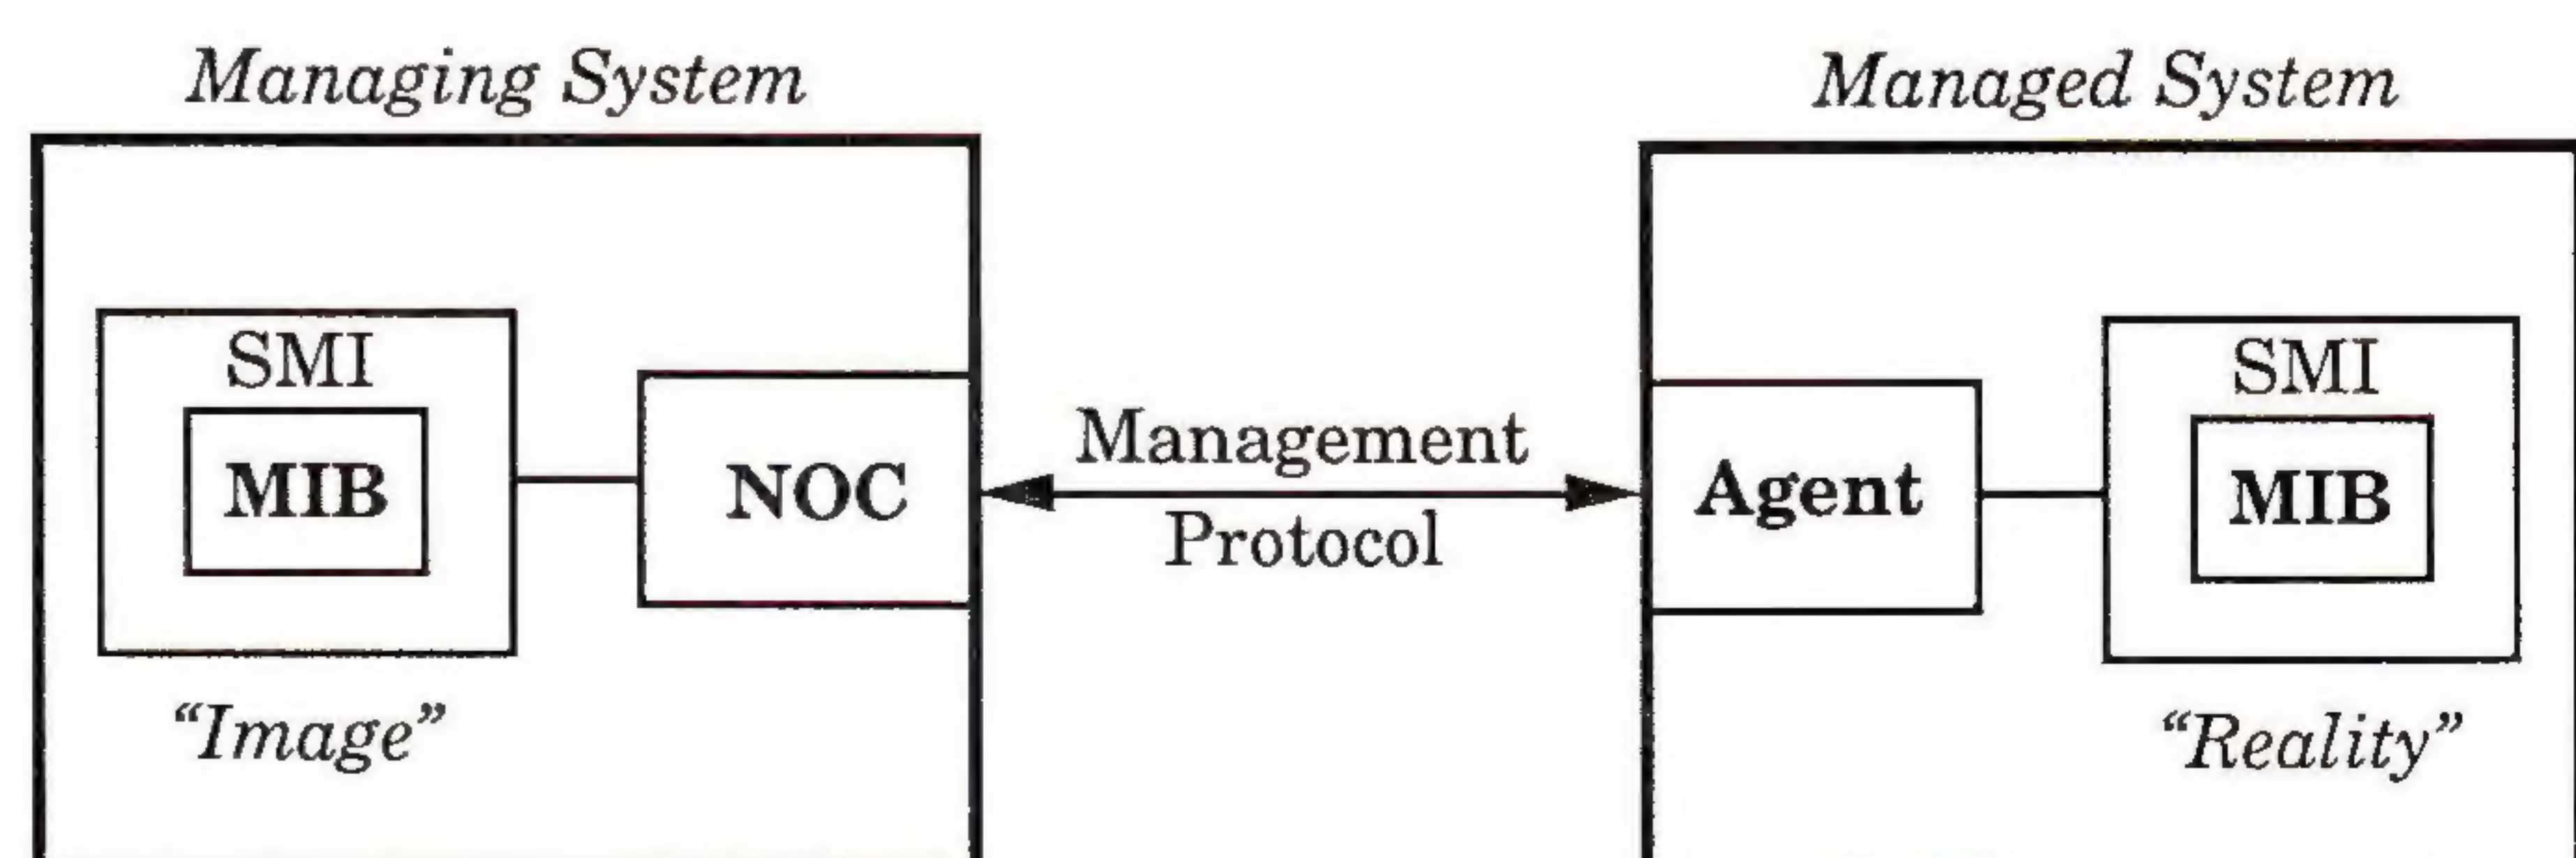


Figure 1: SNMP Management Architecture

continued on next page

MIB Development and Integration (*continued*)

SMI

The *SMI* sets down the elemental structure and form of management information, thereby limiting the universe of choices. Overall, it defines the concept of a MIB as an abstract tree, with individual data items as the leaves. (See Figure 2). It provides a basis for MIB extensibility, in the form of "experimental" and "private" branches to the tree. It establishes *Abstract Syntax Notation One* (ASN.1) as the standard for documenting MIBs [8] and encoding protocol messages [9], picking a small subset of ASN.1's rich supply of data representations. It institutes the basic format for MIB documentation, in the form of an ASN.1 "macro." Finally, it defines the means of identifying individual data items, termed *objects*, for both single instance (scalar) and multiple instance (table) objects, using ASN.1 object identifiers.

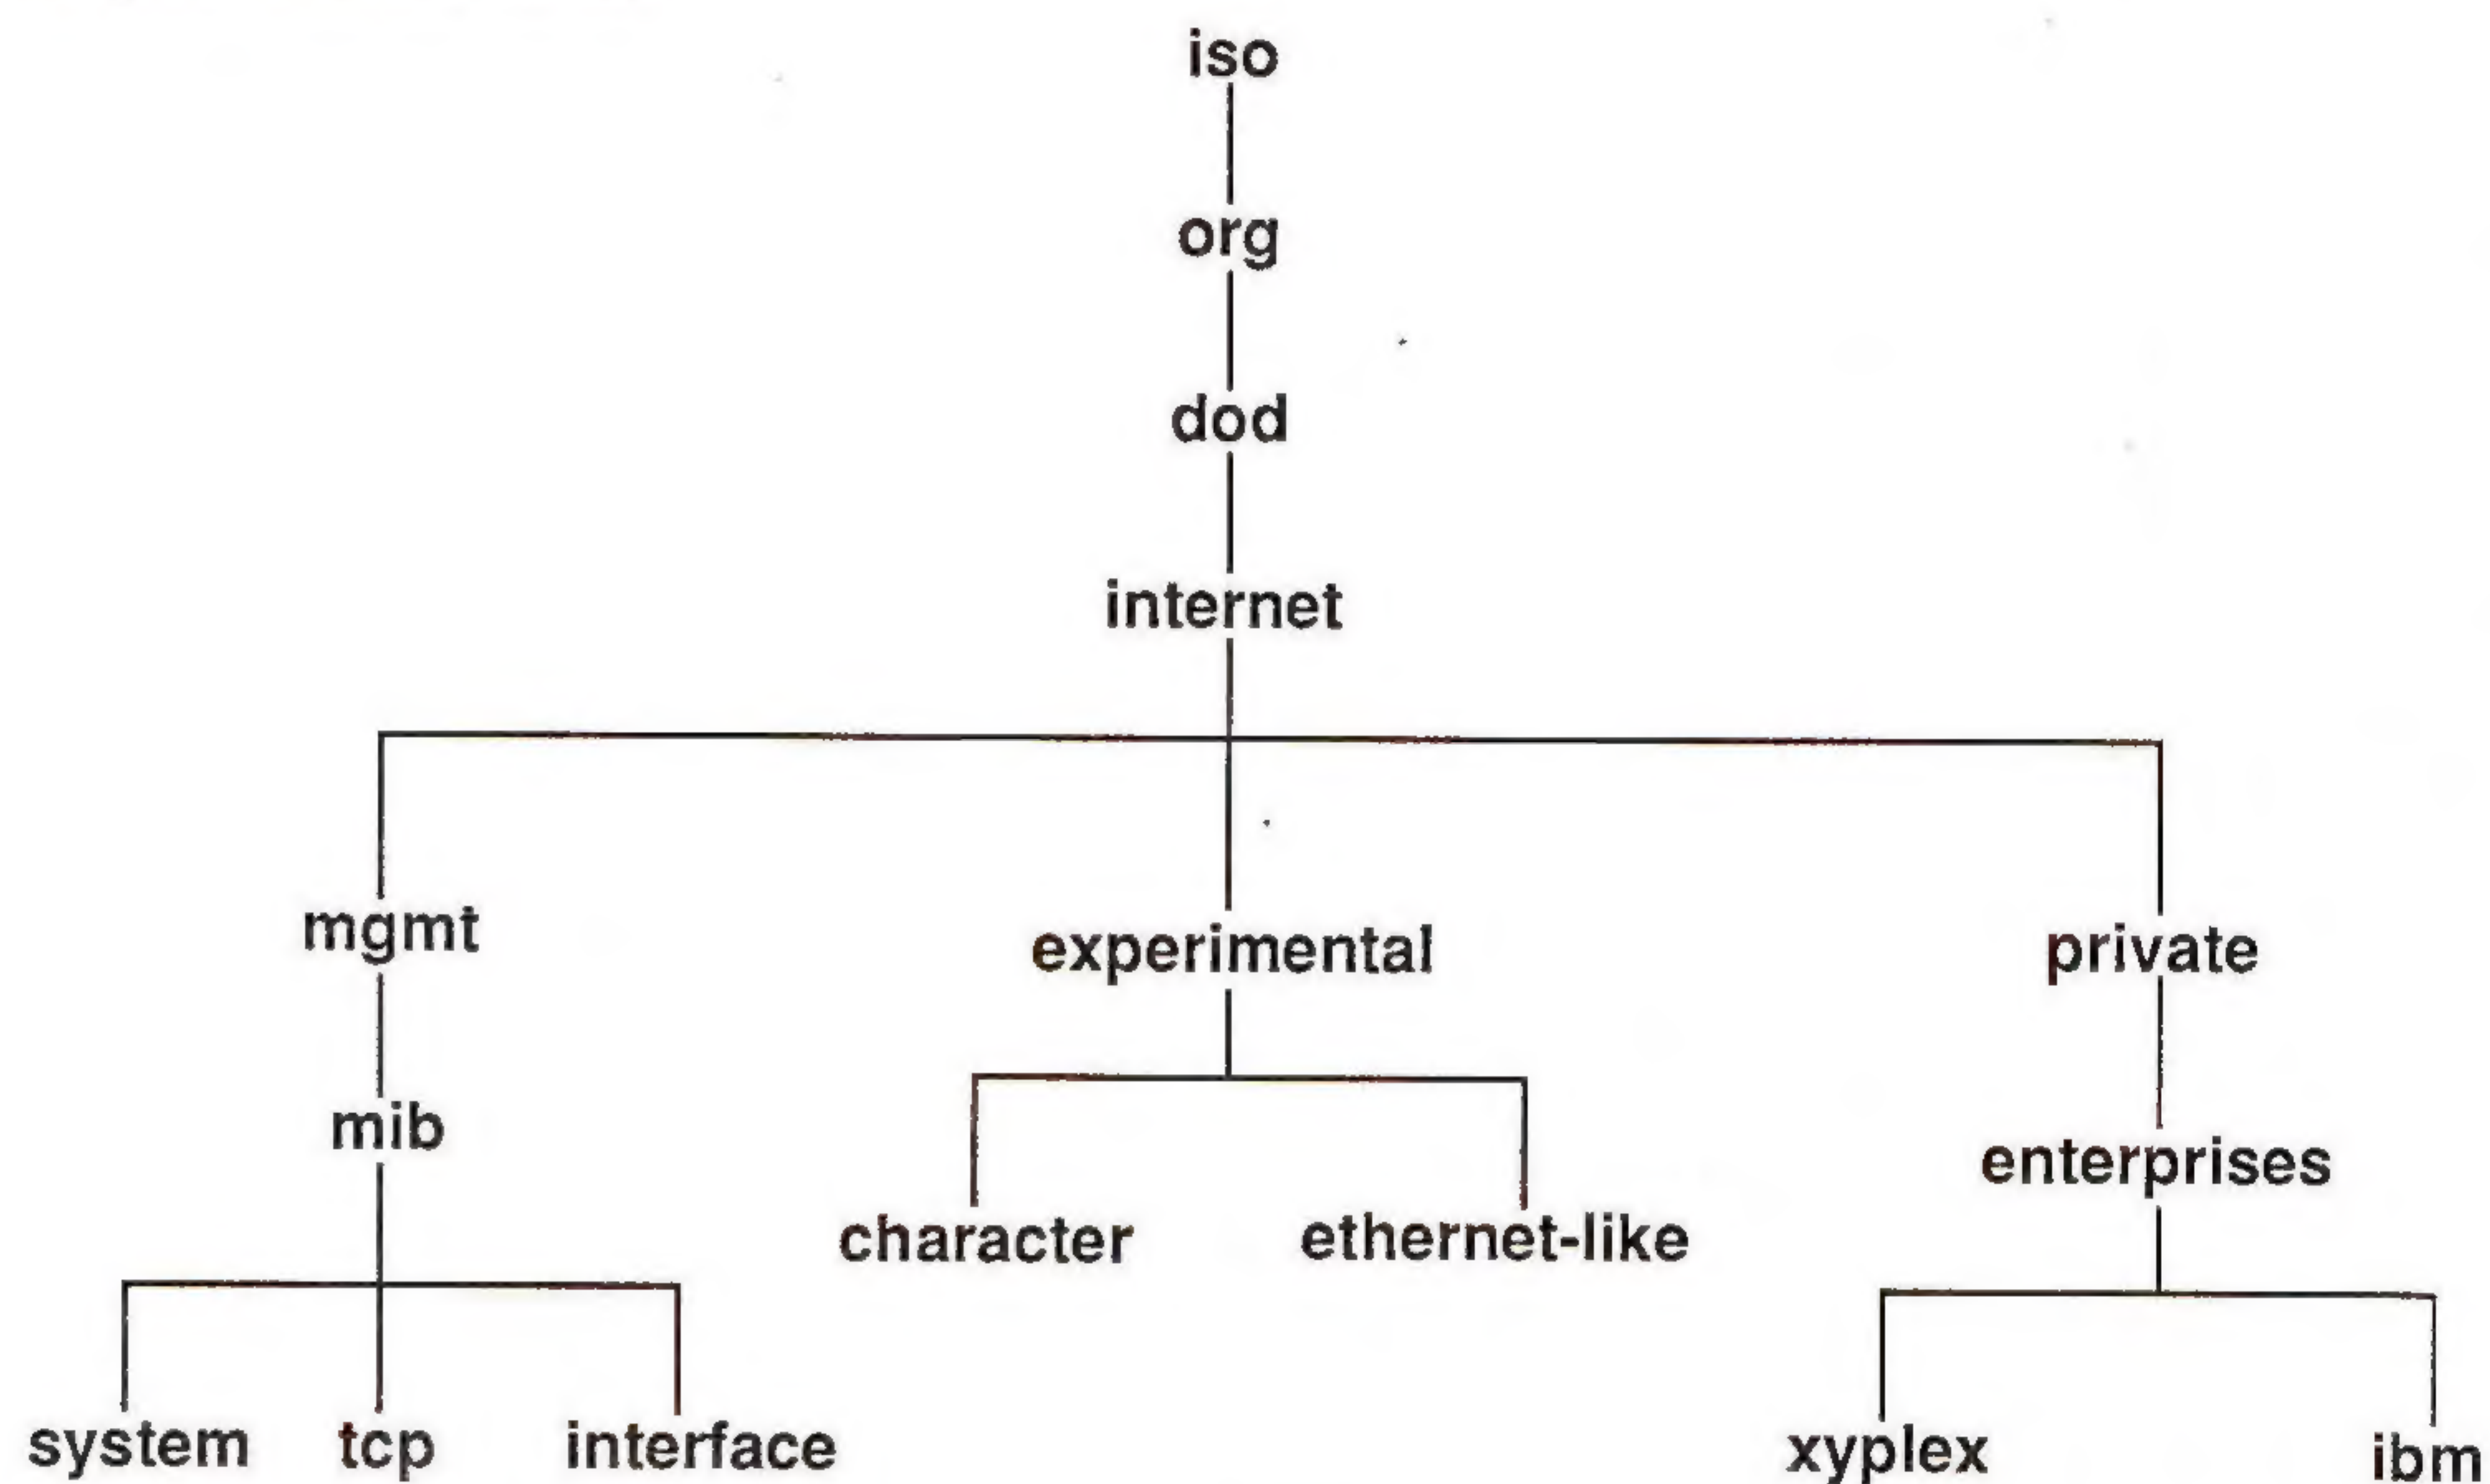


Figure 2: MIB tree

The concept of "tables" deserves special mention, as it has become an area of some controversy. Strictly speaking, tables do not exist in a MIB; the concept was not part of SGMP, and was added for compatibility with CMIP. Abstractly, each data item is an independent entity, in some cases sharing the instance part of their identification. On the other hand, documentation and implementation techniques implied that multiple instance objects exhibit an organization of rows and columns, leading to the desire, and need, to create or delete entire rows. Later documentation [10], attempts to remedy some of the confusion between abstract architecture and practical implementation, but the conflict between the original, pure architecture and the table concept, which has become firmly ensconced in the minds of most MIB designers, remains an area of confusion and awkwardness.

Experience with the SMI proves it as a strong base. It has some readily admitted deficiencies (such as lacking bit strings, integers greater than 32 bits, and text strings distinguishable from binary strings), but as yet the desire for such extensions has not overpowered the desire for continued stability.

MIB

The *MIB* was, in fact, just the starting point for MIB definition. It formed a useful base of management information, and set a tone for future MIB development. The original MIB contains objects for observation and limited control of the Transport, Network, and Data/Physical Link layers of an internet.

Subsequent events show that the MIB designers behaved wisely in severely limiting MIB contents. Although it is easy to find complaints about missing information, the original MIB was interesting in its contents while remaining reasonable to implement, resulting in a sufficient number of implementations to prove the utility of SNMP as a management protocol.

Protocol The *protocol* defines how to exchange messages containing MIB objects, and exactly how to identify instances for the original MIB. The protocol is the level where one chooses either SNMP or CMOT as the carrier for the MIB. As with the SMI and the MIB, the designers kept SNMP simple, preferring a few general purpose functions: *get*, *get-next*, *set*, and *trap*, and specifying it as a request-response protocol using a datagram transport service.

These intentional limitations have considerable impact on MIB design, sometimes keeping the protocol simple at the expense of additional complexity in the MIB. The *get-next* function (often called "the *powerful get-next*") is SNMP's key function, as it gives a NOC the ability to explore a MIB without pre-knowledge of its contents, doing so with a single, relatively simple mechanism in the agent; although MIB designers must take *get-next* into account for efficient access to instanced information.

This is in contrast to CMIP's somewhat more complex mechanisms called *scoping and filtering*. SNMP's *set* function seems simple, but has the painful (for the agent) characteristic that a *set* request can specify many unrelated objects in any order, and the *set* must succeed or fail as a unit. Finally, the *trap* function (for reporting events) is currently the least used and least understood, with considerable controversy over correct event logging operation.

MIB Development The original MIB, now known as *MIB-I*, set the tone for subsequent MIB development, in terms of its form, contents, and the process used to develop it. It originated the following conventions of MIB design:

- *Limited number of objects*: The designers chose a somewhat arbitrary limit of 100 objects for MIB-I. The MIB contains a few more, but the limit imposed discipline, forcing careful evaluation of objects for general utility. This principle leads to the corollary that objects are not duplicated within or across MIBs, and that objects are not defined if the NOC can calculate them more-or-less readily from other objects.
- *Limited computational impact*: The designers accepted a constraint that no more than one counter should appear in critical paths of operation. This constitutes no absolute requirement, but imposes a worthwhile discipline, important in implementing high performance systems. In practice though, it can be difficult to follow, as the definition of "critical path" varies depending on system requirements and implementation models.
- *Limited use of optional objects*: MIB-I contains no optional objects. Instead, entire object groups are implemented if and only if the system has the related capabilities. For example, if the system does not implement TCP, it does not implement the TCP MIB group, but if it does implement TCP, it must implement the complete TCP MIB group. Defining objects as optional allows indecision and compromise on the part of MIB designers with differing opinions, and is at odds with consistency of implementation.

MIB Development and Integration (*continued*)

- *Style of organization and definition:* Conventions in this area are less clear, but MIB-I certainly serves as a model in ways that may not have been intended. For example, its approach to identification of network interface instances with a small, densely-spaced integer that does not change has carried over into various other MIBs, both in reference to those interfaces, and the identification technique used in similar tables. Suffice it to say that the erstwhile MIB designer would do well to understand the reasons behind the organization of the original MIB.

MIB II The second version of the basic MIB, known as *MIB-II* [11], contains incremental changes to MIB-I. Its contributions include:

- *Incremental extensions:* Although the SMI implies that new versions of the MIB were to be new branches in the tree, with retained objects in both the old and new branches, practical experience indicated that the better plan was to make compatible changes to MIB-I, leaving existing objects intact and adding new objects on the ends of the existing branches.

- *Object deprecation:* Some changes required removing old objects, thus MIB-II added the concept of “deprecated” objects to soften the transition to “obsolete” as defined by the SMI. For the convenience of existing NOCs, such objects must be implemented as if “mandatory,” but are subject to removal in a later version of the MIB.

- *Use of textual conventions for object syntax:* Short of adding data types to the SMI, MIB definitions needed a way to indicate common data types, such as an octet string that is to contain only printable ASCII. For this purpose, MIB-II added the use of “textual conventions,” ASN.1 mappings of an object syntax to an existing SMI encoding. For example, `DisplayString` implies an all ASCII octet string for documentation purposes, but the actual ASN.1 encodings for a `DisplayString` and an `OCTET STRING` are the same.

Concise MIB form

During its development cycle, MIB-II became an early example of a new form of MIB definition called *Concise MIB form* [10, 12]. This eliminated much redundancy in MIB definitions, removing opportunities for contradictions and reducing typical document size by over 40%. It established a new convention that techniques for instance identification can appear in the MIB document, rather than the protocol document (for example, SNMP) as the first such did; although instance identification remains architecturally a protocol-specific matter. Also importantly, Concise MIB form extended the ASN.1 macro so more information can readily be parsed by computer programs, further improving the potential for direct integration of new MIBs into existing NOCs with less custom programming.

Private MIBs

The largest body of MIB development is in private MIBs. From the *Internet Assigned Numbers Authority*, organizations can obtain a unique number that defines their private branch of the overall MIB. Within that branch they can do anything they wish, with complete freedom as to process and documentation; although a private MIB should not stray too far from the accepted standards and conventions if it is to have any hope of interoperation with a generic NOC. The typical private MIB defines hundreds of additional objects, and often works only with a NOC customized to match.

In an attempt to help integrate private MIBs into generic NOCs, an individual on one of the SNMP working mailing lists instigated creation of an *Internet MIB Repository*. The result is a public collection of MIB documents, available on the Internet at a neutral, central site, documenting the private MIBs of any organization wishing to submit a document that will pass scrutiny by a MIB syntax checking program (a MIB compiler). In particular, system vendors who implement private SNMP MIBs can now easily supply documentation that can be processed by a growing number of NOCs.

Experimental MIBs

The experimental branch of the overall MIB contains a growing number of experimental MIBs. As with private MIBs, each has its own branch of the tree, obtained from the Internet assigned numbers authority. Such MIBs often are developed by Internet working groups under the auspices of the *Internet Engineering Task Force* (IETF). The section on MIB process, following, describes such development in more detail. Experimental MIBs are subject to considerable discussion by interested experts and represent a common agreement that is assumed subject to change based on experience in implementation and use. They can thus be widely implemented in a common way with the understood risk that they will be replaced, nicely blending interoperability with the flexibility to learn and improve.

Some examples of experimental MIBs include various flavors of IEEE 802 devices, including Ethernet/802.3, Token Ring/802.5, and bridges. Other experimental MIBs cover the proposed Point-to-Point Protocol (PPP), and character stream devices. As experimental MIBs mature, they may move to the management branch and become Proposed Standards.

MIB Process

For discussion purposes, the following steps describe a development process for a new standard MIB branch. The process for a private MIB would be less formal, and typically internal to its developing organization. The process described here does not reflect documented, official IAB or Internet procedure, as such does not yet exist; instead, the following is based on the author's ongoing experience in developing a MIB.

Establish working group

Internet standards development is the task of the IETF. The IETF is divided into Areas, such as Applications, User Services, and Network Management, each with an *Area Director*. Area Directors oversee *Working Groups* that exist for the duration of specific tasks, such as development of an experimental MIB. Anyone who can demonstrate sufficient interest can cooperate with the appropriate Area Director to establish a working group. Subject to approval by the Area Director, a working group writes its own charter, which describes its task and sets milestones to plan and measure progress.

For a MIB extension, the working group should attract the participation of multiple parties with expertise and interest in the subject area, usually representing an assortment of businesses and educational institutions. Indicative of Internet community spirit, and the strong commercial interest in Internet protocols, such a working group typically exhibits remarkable cooperation among businesses that otherwise are ardent competitors.

Develop Draft

The major part of a MIB working group's task is to develop an *Internet Draft* defining the MIB. Typically the draft will have one or two editors, with starting technical contributions from several group members. As each draft appears, the group reviews it, comes to sufficient consensus on issues, and the editor issues a new draft.

continued on next page

MIB Development and Integration (*continued*)

Once the group has general consensus, the document becomes an Internet Draft, publicized and readily available within the IETF, supporting broader review and comments by IETF members who did not participate in the working group. The draft may position the MIB extension in the management branch of the MIB or, more commonly for development purposes, in the experimental branch.

Most of a working group's task is coordinated and carried out via electronic mail on the Internet. Final decisions are usually made at open IETF meetings, held three to four times a year. Although it is possible to participate in a working group without electronic access to the Internet, such participation would be severely handicapped.

For a MIB, the working group must find those objects that are of general interest and use to network managers without excessive impact on implementations, either in complexity or efficiency. Object design must avoid assuming particular implementation models; instead design must proceed from a sound architectural base. Experience with existing implementations or standards provides important input, but cannot be followed blindly, as design requirements and goals may differ.

To help assure acceptance, the MIB must be pared to the bone. A big, complex MIB will not enjoy the widespread implementation necessary to prove its utility. The working group will have to agonize over inclusion of each object, often omitting objects that would be interesting but are too costly or not sufficiently representative of differing implementations. Even so, such objects are not lost, as they can easily be included in private MIBs. Also, the group must be aware of other MIBs that exist or are in development, so as to avoid duplication.

The working group must anticipate how information is to be used, especially with tabular information. For example, table access may require searching for individual entries or sequential access to all entries. In the former case, SNMP's *get* and *get-next* functions work best with an instance identification that directly selects the desired instance, such as the destination internet address of a routing table entry. In the latter case, *get-next* is the likely function of choice, and a direct identification may force the agent into a complex search for each entry as the NOC requests it. In this case, the best instance identifier may be a simple, sequential integer that bears no direct relationship to other instance parameters. If both types of access are anticipated, the MIB may need both organizations, or the working group may have to make a hard choice.

Implement

Once the working group has a stable Internet Draft, the MIB must be implemented. The first implementation does not have to be a carefully designed commercial product. It may be a prototype, sufficiently functional for interaction with a NOC or two, but not intended for open public use. Its purpose is to prove that the MIB can be reasonably implemented and has none of the flaws that only become obvious when a programmer attempts committing abstract architecture to hard code. Clearly, if implementation uncovers problems, the working group must reiterate enough of the process to fix the shortcomings, and the corrections must in turn be implemented.

Publish RFC

The IAB decides when an Internet Draft is to enter the Internet standards track as a *Proposed Standard* in the form of a *Request For Comments* (RFC). Several steps lead to this point. Given a preliminary implementation and resolution of all outstanding comments, the working group presents their explicit recommendation to the Area Director. The Area Director is responsible for first approval, then passes the package to the *Internet Engineering Steering Group* (IESG). A positive consensus in the IESG results in submission to the IAB for final approval. Upon approval by the IAB, the IESG Secretary gives the RFC text to the RFC Editor for publishing. At this point, several experimental MIBs have also been moved into the management branch.

Implement widely and use

Once available as an RFC the MIB must be widely implemented and used in real-life network management. If this does not occur and the causes of such lack of interest are not fixed, the MIB is dead. If this does occur, the experience gained will indicate any changes that must occur before the next step, from *Proposed Standard* to *Draft Standard*. If the necessary changes are extensive, the MIB goes back to the working group and Internet Draft status.

Declare Standard

When a MIB has survived the test of widespread implementation and use, the IAB can promote it to *Draft Standard* and eventually to *Internet Standard*, with additional status of *Recommended* or *Required*.

Protocol future

SNMP and the SMI will change slowly, if at all. Although pressure exists now to consider changes to both, the desire for stability and deeper experience with what we already have continues to override the desire to fix even the admitted problems, much less open the door for change in more controversial areas.

The SNMP MIB will expand. New MIBs, in the management, experimental, and private branches, continue to appear at an astounding rate. The rate of MIB invention outstrips the ability of the management community and the NOC developers to employ the MIBs for the most benefit of network managers. Instead, too many NOCs concentrate on a few functions and superficially impressive graphics, leaving the hapless network manager with relatively primitive means to use even the standard MIB, on an object by object basis. Most needed is a better way to integrate new MIBs into existing NOCs. The following section on interoperation addresses this.

CMIP ?

Finally, what of OSI and CMIP? Networking vendors continue to express support for CMIP while implementing SNMP. Given its weight as an International Standard, use of CMIP seems inevitable, but timing remains unclear. As with SNMP, the biggest problems facing CMIP are MIB definition and widespread implementation in managed systems. MIB definition is hard in the friendly, informal Internet standards environment, and that difficulty increases in the more political OSI environment. Furthermore, it may be that CMIP's additional complexity could delay widespread implementation indefinitely. The SNMP community continues to develop new MIB objects although many of them remain unimplemented, and this is in the face of a design discipline not imposed on CMIP, thus CMIP's problems can only be worse.

Part of the answer for practical OSI network management may lie in incorporation of some SNMP concepts, but, as yet, no one from the SNMP community has found the time and energy to take SNMP into the international standards process.

continued on next page

MIB Development and Integration (*continued*)

Interoperation

The MIB explosion must continue to make necessary information available, but NOC and agent suppliers need a way to reconcile their products. NOC suppliers cannot track every MIB development, and agent suppliers cannot do or expect custom programming for every NOC. The answer to this dilemma may be implied by the current practices, that is, standard, machine-readable MIB descriptions have proven to be valuable resources for minimal integration of a new MIB into an existing NOC. A growing number of NOC vendors supply a means to directly read such information.

Although the current MIB format can supply a computer program with most of the information needed to interoperate, it falls far short of the information needed to provide useful monitoring and control to a network manager. This does not imply that such information should be added to the basic form. To do so would adversely impact the time needed to get a new MIB implemented. Rather such information should be available from supplemental documents. The following proposal is still in embryonic form. It has enjoyed a bit of discussion, but does not represent work in progress. The supplemental information is in two categories:

- *Common Supplement:* The Common Supplement provides information applicable across implementations. Its primary purpose is to assist in formatting displays. It has clear, full-word labels for objects, groups objects logically, both within a display and into different displays, and provides help text aimed at network managers rather than MIB implementors. It defines how to create table entries, and indicates what objects in such an entry are required. Further, it might go into analysis of how to recognize problems and what to do next when they occur, although attempting this level of utility could make the idea too complex for practical realization.

- *Implementation Supplement:* The Implementation Supplement adds implementation-specific information to the common supplement. It further refines the Common Supplement, possibly adding additional labels or groupings, and it indicates what objects are actually implemented for those cases where objects are optional, or for non-standard implementations that do not include mandatory objects. It provides a means for the vendor to document the behavior of a particular agent implementation, in a form that should be shipped as part of the product.

Summary

We have examined the origin, philosophy, and structure of SNMP to support understanding of the MIB and its place, both in design and in use. Our examination of MIB evolution lead to a detailed description of the process for developing a MIB extension. Finally, after predictions of protocol futures, and recognition of the growing problem of new MIBs, we proposed additional standard-format, machine-readable documentation as an approach to improving interoperability.

Acknowledgements

In the little more than a year I've been involved with SNMP, I owe much to Jeff Case, Chuck Davin, Keith McCloghrie, Marshall Rose, and numerous other individuals on the Internet for their help and patience in my education regarding SNMP.

References

- [1] V. Cerf, "IAB Recommendations for the Development of Internet Network Management Standards," RFC 1052, April 1988.
- [2] V. Cerf, "Report of the Second Ad Hoc Network Management Review Group," RFC 1109, August 1989.

- [3] J. Davin, J. Case, M. Fedor, & M. Schoffstall, "A Simple Gateway Monitoring Protocol," RFC 1028, April 1989.
- [4] U. Warrior & L. Besaw, "The Common Management Information Services and Protocol over TCP/IP," RFC 1095, April 1989.
- [5] M. T. Rose & K. McCloghrie, "Structure and Identification of Management Information for TCP/IP-based internets," RFC 1155, May 1990.
- [6] K. McCloghrie & M. T. Rose, "Management Information Base for Network Management of TCP/IP-based Internets," RFC 1156, May 1990.
- [7] J. D. Case, M. S. Fedor, M. L. Schoffstall, & J. R. Davin, "Simple Network Management Protocol," RFC 1157, May 1990.
- [8] Information Processing Systems, Open Systems Interconnection—Specification of Abstract Syntax Notation One (ASN.1), ISO, International Standard 8824, December 1987.
- [9] Information Processing Systems, Open Systems Interconnection—Specification of Basic Encoding Rules for Abstract Syntax Notation One (ASN.1), ISO, International Standard 8825, December 1987.
- [10] M. T. Rose, K. McCloghrie (editors), "Towards Concise MIB Definitions," Internet Draft, IETF, December 1990.
- [11] "Information Base for Network Management of TCP/IP-based Internets: MIB-II," RFC 1158, May 1990.
- [12] "Information Base for Network Management of TCP/IP-based internets," Internet Draft, IETF, December 1990.
- [13] *ConneXions*, Volume 3, No. 3, March 1989, Special issue on Network Management.
- [14] *ConneXions*, Volume 4, No. 8, August 1990, Special issue on Network Management and Network Security.
- [15] Marshall T. Rose, "The Simple Book—An Introduction to Management of TCP/IP-based internets," Prentice-Hall, 1990, ISBN 0-13-812611-9.
- [16] Paul Brusil, "Components of OSI: Systems Management," *ConneXions*, Volume 5, No. 4, April 1991.
- [17] Lee LaBarre, "Management By Exception: OSI Event Generation, Reporting, and Logging," Invited Paper in Proceedings of the IFIP *Second International Symposium on Integrated Network Management*, Washington, D.C., North-Holland Publisher, April 1991.

[Ed. This article is based on a paper presented at the 1991 Silicon Valley Networking Conference. Printed with permission].

BOB STEWART has been designing and implementing network and system management software since 1969. He was the DECnet network management architect from 1976 to 1982, took a break to write the initial architectural specification for the IEEE 802 MAC bridge, then barely kept his shirt after helping found a home computer software company in 1984. Since then he has designed and implemented software for Xyplex, where his title is Network Architect and his name plate says *Architectural Theologian*. He is active in the IETF SNMP working group, and is chairman of the Character MIB and Special-purpose Hosts working groups. He collects guns, swords, and puns, and has opinions on almost everything. He can be reached as: rlstewart@eng.xyplex.com.

SNMP Security

by

Keith McCloghrie, James R. Davin and James M. Galvin

Background

The *Simple Network Management Protocol* (SNMP) [2] was an outgrowth of SGMP, the Simple Gateway Monitoring Protocol [1]. SGMP was a grass roots effort to define a network management protocol to meet the need to monitor IP gateways, a need which in 1987 had become critical for maintaining the Internet. Because of the immediate need, the SGMP was designed with the minimum set of features required. Among the features not considered immediately required were the ability to control gateways and security. Without full support in the SGMP for the control of gateways, the lack of security was considered acceptable. However, a hook was included to enable the later addition of authentication.

Both because of the SGMP's ability to fulfill the narrow focus of its goals, and because of the lack of an alternative, SGMP was successful. So successful in fact, that the recommendation of the first *ad hoc* committee meeting to set direction for network management protocols in the Internet [3], was that SNMP, as a successor to SGMP, should be the "short-term" standard. The subsequent definition of SNMP included support for controlling as well as monitoring any/all network devices, not only IP gateways. However, security still gained no more than another hook for its later addition. This hook used a "community" field to identify the administrative relationship between the sender and the receiver, including the authentication algorithm being used. For the time being, only the "trivial" authentication algorithm was defined, in which any community name known to the receiver was automatically authentic. [Note that because the community field identifies, among other things, an algorithm, it is more than just a "password."]

Subsequently, several members of the SNMP community began thinking about how to define additional authentication algorithms to provide a reasonable level of security protection. With the growth of SNMP such that it has now become the *de facto* standard for management of not just TCP/IP networks, but increasingly of other networking regimes also, the lack for security is fast becoming a problem. At last (two years later), the effort to define additional algorithms has now resulted in the publication of Internet Draft documents [7, 8, 9], ready for trial implementations. The length of time it has taken is an indication of how hard security can be, especially for network management which has the need to continue operating under adverse network conditions.

The threats

As in all security work, the threats that may be encountered must be identified. For SNMP Security, the identified threats are:

- *Modification of Information:* The threat that an in-transit message from an authorized source may be modified. For example, a *Set-Request* to further tighten the filtering in a bridge could be modified to specify no filtering.
- *Masquerade:* The threat that management operations not authorized for some source may be attempted by that source by assuming the identity of another source that has the appropriate authorizations. For example, any network user pretending to be the NOC's management station.

- *Message Stream Modification*: The threat that messages may be re-ordered, delayed or replayed to cause unauthorized management operations to be performed. For example, capturing an authorized request to reboot a router, for replay at a later date.

- *Disclosure*: The threat of eavesdropping on the exchanges between managed agents and a management station. Protection against this threat is mandatory when SNMP is used to administer private parameters on which its security is based. It would also be appropriate when setting passwords in, say, a terminal server.

Goals and constraints

To protect against these threats, SNMP Security must provide: message integrity, data origin authentication, replay protection, and privacy.

In addition, the design was influenced by the following constraints:

- When requirements of effective management in times of network stress are inconsistent with those of security, the former are preferred.
- Neither SNMP nor its underlying security mechanisms should depend upon the ready availability of other network services (e.g., *Network Time Protocol* (NTP) or secret/key management protocols).
- A security mechanism should entail no changes to the basic SNMP network management philosophy.

Mechanisms

Three basic mechanisms are used to meet the goals: use of the *MD4* message digest algorithm, use of the DES cryptographic algorithm, and use of loosely synchronized clocks.

The MD4 [4] message digest algorithm is used to support message integrity. The MD4 calculation is performed on the concatenation of a secret value and a portion of an SNMP message, and the resultant 128-bit digest is included as part of the message sent to the recipient. Thus, a modified message is only valid if the digest it contains is correspondingly modified, but the correct digest cannot be calculated without knowing the secret (which, of course, is not transmitted as part of the message).

Data origin authentication is provided by the same MD4 digest calculation, since if only the originator of the message and its intended recipient know the secret value used in the calculation, and if the recipient did not send it, then it must have been generated by the originator.

Replay protection is provided by including in each message a timestamp value indicating the time of the message's generation according to the clock maintained by the originator. The recipient uses the timestamp to determine both that the message is recent, and that it was generated subsequently to all previous messages received.

The *Data Encryption Standard* (DES) [5] in the Cipher Block Chaining mode of operation [6] is used to provide privacy. An appropriate portion of the message is encrypted prior to being transmitted to its recipient.

Problems of integrating the mechanisms

Identifying these basic mechanisms was the easy part of the task. What turned out to be much harder was resolving the issues of how to make them usable by SNMP implementations.

continued on next page

SNMP Security (*continued*)

First, consider that data origin authentication and message integrity are based on the *source* of the message, but encryption is based on the *destination* of the message. Access control needs to be based on which source is trying to access what target, and the target is based on the destination. Thus, there is a need to differentiate between source and destination, a need which the existing “community” field does not provide. Second, how are the loosely synchronized clocks maintained, especially if they cannot be based on NTP? Third, how are the secrets distributed and maintained?

Introducing the SNMP Party

After much consideration, it was decided that SNMP’s Administrative Framework needed to be refined to be able to distinguish between source and destination. This has been done with the introduction of the concept of an SNMP “Party.” An SNMP party is defined as an execution context of a SNMP protocol implementation. Whenever a SNMP protocol implementation processes a message, it does so by acting in the role of one of the SNMP parties defined for it. Each SNMP party executes at a specific transport address, and has specific authentication parameters, privacy parameters, proxy information, and a MIB view. The authentication parameters include an algorithm, a secret, and the state information needed to maintain its clock and ensure proper message ordering. The privacy parameters include an algorithm and a secret. The proxy information either indicates no-proxy, or “points” to another SNMP party where the real-agent executes. The MIB view specifies the subset of an agent’s MIB that the party can access.

By including an algorithm in both the authentication and privacy parameters, multiple parties with different capabilities can be defined for execution at a single SNMP protocol entity. One party can use no-authentication and no-privacy; another can use MD4-authentication and no-privacy; and another can use MD4-authentication and DES-privacy. However, for security reasons, the use of privacy requires the use of authentication.

In this model a SNMP message is originated by one party and destined for another party. The message is authenticated according to the authentication information of the source party. The message is encrypted (or not) according to the privacy information of the destination party. Access control specifies that a specific source party is allowed to originate a particular set of SNMP operations (e.g., *Get-Request* and *Set-Request*, or just *Get-Request*) to a specific destination party. In practice, each SNMP protocol implementation needs to keep a local database of party information, both for parties that execute locally and for remote parties with which local parties communicate.

While this model requires the format of a SNMP message to change in order that both the destination and source parties and the appropriate authentication information (e.g., the digest and the timestamp) can be sent with the message, it does not (fortunately) entail any change to the SNMP PDU. In fact, an agent implementation which followed the guidelines in the original SNMP protocol specification, should be able to implement the new message format with little more than changes to the authentication service function.

Maintaining clocks

Each SNMP party has a (relative) clock. In order to be authentic, a message received by the party must have a timestamp which, when added to an administratively defined “lifetime,” is greater or equal to the value of the clock at the time of receipt.

Thus, in order to communicate with a remote party, a SNMP protocol entity must retain in its local database a clock value which is (loosely) synchronized with the remote value. Since clocks tend to drift, it is necessary for the clocks to be inspected periodically, and re-synchronized if necessary. This chore is delegated to the management station. However, a few features are included in the protocol to enhance the synchronization through the regular exchange of messages, and the party clocks are purposely positioned in the MIB to support easy read access to them through unauthenticated and authenticated *Get-Requests*, so that a manager does not need to maintain synchronization with all its agents all of the time, but can let some lapse, and later re-synchronize when necessary.

The lifetime value must be carefully chosen. It must be large enough to accommodate variations in communications delay as well as to accommodate a small amount of drift. On the other hand, it is the lifetime value which provides the window in time during which a message is valid, and so lifetimes must be kept within the bounds which an administration wants protection against replay attacks (e.g., a few minutes).

The constraint that network management must continue to operate even under conditions of network stress, even if this impacts the level of security, was mentioned above. One way this can be achieved, is by a manager artificially advancing its notion of the party clock in an agent, so that even though communication delays may have increased dramatically, a message will still be considered authentic when it arrives at an agent.

Distributing secrets

The use of both the MD4 authentication and the DES privacy algorithms rely on secrets, which are shared by the originator and the recipient. If these algorithms are to maintain their level of security, their secrets must remain secret and not be available to would-be attackers. Thus, they cannot be transmitted over the network, except in an encrypted form. So, once the two SNMP protocol entities share a secret used for encryption, that secret can be used to encrypt new (changed) secret values for that or any other party. Such changing of secrets on a regular basis is very desirable from a security standpoint.

However, the initial distribution of secrets cannot be done over the network. Instead, it must be done manually as a piece of initial configuration information entered into the manager and the agent, before secure communication is possible. Subsequent distribution of secrets can be done via SNMP access to appropriately secured MIB objects, unless one or both entities loses its knowledge of the secret (or if it is known that the secret has been comprised).

Identifying SNMP Parties

Since each SNMP party is unique to the particular SNMP protocol implementation where it executes, many parties need to be defined. A convenient way to do this is to identify them by *Object Identifiers* (OIDs), of which there is an infinite supply! This allows each network administration to obtain its own branch in the OID tree, and allocate OIDs for its agents from there.

However, to simplify matters, a set of six "initial" OIDs have been assigned for use with each IP address, three for local execution at an agent, and three for the agent to communicate with. The three have different settings of authentication and privacy algorithms, with an appropriate MIB view and access control parameters defined for each. The extension of these six to the number actually required in an agent can, of course, be done through the use of SNMP requests acting on appropriate MIB objects.

continued on next page

SNMP Security (continued)

Summary

Security of management operations is a feature which has been overdue for inclusion in SNMP. Due to the operational necessity, use of SNMP for monitoring has flourished even without security. However, some vendors and network administrators have been slow to use SNMP for controlling their devices until security is in place. Now that specifications of the SNMP Security Protocol outlined above are available for experimental implementations, we will be able to determine how practical these specifications are, and (hopefully) can look forward to secure network management becoming a reality in the not-too-distant future.

References

- [1] J. R. Davin, J. D. Case, M. S. Fedor, & M. L. Schoffstall, "A Simple Gateway Monitoring Protocol," RFC 1028, November 1987.
- [2] J. D. Case, M. S. Fedor, M. L. Schoffstall, & J. R. Davin, "Simple Network Management Protocol," RFC 1157, May 1990.
- [3] V. Cerf, "IAB Recommendations for the Development of Internet Network Management Standards," RFC 1052, April 1988.
- [4] Ronald L. Rivest, "The MD4 message digest algorithm," RFC 1186, October 1990.
- [5] FIPS Publication 46-1, "Data Encryption Standard," National Institute of Standards and Technology (NIST), US Department of Commerce, January 1977.
- [6] FIPS Publication 81, "DES Modes of Operation," NIST, Dec. 1980.
- [7] J. R. Davin, K. McCloghrie, J. M. Galvin, "SNMP Administrative Model," RFC in preparation.
- [8] J. M. Galvin, K. McCloghrie, J. R. Davin, "SNMP Security Protocol," RFC in preparation.
- [9] K. McCloghrie, J. R. Davin, J. M. Galvin, "Experimental Definitions of Managed Objects for Administration of SNMP Parties," RFC in preparation.

KEITH McCLOGHRIE is an Associate Director of Engineering at Hughes LAN Systems, Inc. where he is responsible for the development of network management products. He is a member of the IETF's Network Management Directorate and has been an active member of the SNMP working group since its inception, involved in the development of many MIB specifications. He is a member of the IFIP Working Group 6.6 on network management, involved in the organization of the International Symposia on Integrated Network Management. He gained his B.Sc. in Mathematics from Manchester University in England.

JAMES R. DAVIN currently works in the Advanced Network Architecture group at the M.I.T. Laboratory for Computer Science where his recent interests center on protocol architecture and congestion control. He serves on the steering group of the IETF where he contributes to the evolution of standards for network management. In the past, he has been engaged in router development at Proteon, Inc., where much of his work focused on network management. He has also worked at Data General's Research Triangle Park facility on a variety of communications protocols. He holds the B.A. from Haverford College and masters degrees in Computer Science and English from Duke University.

JAMES M. GALVIN is a Senior COMSEC Scientist at Trusted Information Systems, Inc. in Glenwood, Maryland. Dr. Galvin's responsibilities emphasize communications security, especially computer networks, architectures, policies, and procedures. He is a principal in the development of TIS' soon to be openly available implementation of Privacy Enhanced Mail. He is very active in the IETF Security Area Advisory Group and Chair of the OSI Implementor's Workshop Security Special Interest Group, hosted quarterly by the NIST. He received his Ph.D. and M.S. degrees, both in Computer Science, from the University of Delaware in 1988 and 1986, respectively. In 1982, he received his B.S. in Computer Science and Mathematics from Moravian College in Bethlehem, PA.

The NREN Takes Shape

by Mike Roberts, EDUCOM

- Introduction** The announcement in February of the federal *High Performance Computing and Communications* program (HPCC) marked a major turning point for advocates of the *National Research and Education Network* (NREN). After five years of study and planning, the Bush Administration has embraced the concept of an advanced computer network to tie together research and education activities in the United States.
- The program calls for an interagency effort to be directed by the *Federal Coordinating Council on Science, Engineering and Technology* (FCCSET), which reports to Dr. Allan Bromley, the Director of the White House *Office of Science and Technology Policy* (OSTP). The announcement says, in part, "High performance computing and computer communications networks are becoming increasingly important to scientific advancement, economic competition and national security. The technology is reaching the point of having a transforming effect on our society, industries and educational institutions. The goal of the HPCC program is to accelerate significantly the commercial availability and utilization of the next generation of high performance computers and networks..."
- Funding** A total of \$638 million in federal funding for HPCC is proposed in the President's FY92 budget, of which \$92 million would go to the NREN. The balance is distributed among categories of high performance computing systems, advanced software technology, and basic research and human resources. Within the NREN category, major chunks of money go to the *Defense Advanced Research Projects Agency* (DARPA) for research and development in gigabit technology (\$32.9M), to NSF for operational costs of the federal portion of the NREN (\$32.7M), to the Department of Energy for network applications research (\$12M), and lesser amounts to several other agencies. Of the total \$638 million in HPCC funds, \$149 million are characterized as new for FY92.
- Gore Initiative** On a parallel but separate track from that of the Administration, Congress has been considering high performance computing and networking legislation for the last two sessions. Spearheaded by Senator Albert Gore, Jr (D-Tenn), bills have been drafted in both houses which establish a statutory legal basis for the NREN, thereby providing some insurance from the yearly agency battles over budget priorities. To some extent, the separate Congressional and Executive Branch programs reflect party politics between Democrats and Republicans, with each side seeking advantage for the next election battle. In recent hearings, progress has been made in bringing the two different approaches into greater harmony, and there is the prospect that the Administration will support the NREN bills when they come to the floor for action later in the year.
- Effects** The NREN program will have effects on the Internet at several levels, almost all of them positive. Perhaps most important, the existence of the program, and the likely passage of legislation embedding the rationale for the network in federal law, will create a high level priority for advanced networks and network technology within the federal government's overall research and development budget, a Mississippi River of funding which will reach nearly \$80 billion in FY92. In the wacky world of Washington science politics, this credentialing of the program is essential.

continued on next page

The NREN Takes Shape (*continued*)

Direction of the NREN program by OSTP/FCCSET will also provide a unified focus for federal networking efforts for the first time in many years. Not since the early Defense Department sponsorship of the ARPANET [4] has there been a single point for coordination of networking strategy and funding.

New federal research networking funds, estimated at \$100 million a year over the next several years, will serve an important role in leveraging additional network investment within the research and education community. In the last several years, university and state level network capital expenditures have been in the range of several hundred million dollars, or nearly ten times the level of federal funding.

A forcing function

The NREN is also serving as a forcing function within the Internet community to address some pressing structural problems, including the issues of standards development, domestic versus international growth of the Internet, and federally supported versus commercial access to and use of the network. The *Federal Networking Council* (FNC) and its newly formed external advisory group, the *Federal Networking Advisory Committee* (FNAC), have established several working groups, in coordination with the IAB/IETF, to address these and other policy matters. [1, 2, 3, 5].

Debate on the NREN and advanced networking technology in general has also begun to open up broader consideration of national telecommunications policy, a highly politicized arena where high stakes special interest wars have been fought for more than ten years. Several recent forums have urged that the federal government take the lead in forging an advanced communications infrastructure for the entire country based on broadband fiber optic networks.

Membership

When it officially comes into existence next October at the beginning of the federal fiscal year, the NREN will be composed of parts of existing federally sponsored and supported research and education networks. These include all of NSFNET [6, 7] and major portions of networks of other federal departments with research programs, including NASA, Energy, DARPA, NIH, NOAA, etc. As is the case with the Internet today, boundary definitions for the NREN will be hard to come by. Some network users and proprietors will no doubt vote themselves into membership in the NREN. There are at least two defining conditions for membership—a commitment to use of the network to advance research and education, and a commitment to development and use of high performance networking technology—minimally gigabit speeds—to support the goals of the network within research and education.

Expansion

There is considerable interest within the networking community concerning the deployment of gigabit technology in the NREN, as well as in the Internet. The official game plan, as spelled out in FNC planning documents and draft Congressional legislation, is that the NREN program will be carried out in two dimensions—connectivity and bandwidth—simultaneously. The National Science Foundation (NSF) is assigned the responsibility, on behalf of FCCSET, to expand the operational reach of the NREN so that it connects universities, colleges, libraries to research sites and to supercomputer centers and information services providers. Over the next several years, the number of connected sites and campuses will likely double from the present six or seven hundred to as many as fifteen hundred.

At the same time, DARPA is assigned the R&D role, building on the work it is already doing with the gigabit trials and testbeds. To design, develop and deploy gigabit technology in the NREN is a large undertaking, especially since the goal is not just to have network links with gigabit bandwidth, but to provide gigabit bandwidth end to end between advanced workstations and supercomputers.

A target

There has been a tendency in Washington to view the NREN as a thing and as an end in itself. Neither is true. The NREN is not a single network, but part of the family of Internet networks. The performance goal is to recreate the situation which existed in the 1970s in which American networking technology was at the international leading edge of both R&D and operational networking. The draft legislation establishes a target of usable gigabit bandwidth on the NREN by 1996. When we reach the target, there will be new challenges in bandwidth, connectivity, and applications.


References

- [1] Postel, J., "An overview of the Internet Activities Board," *ConneXions*, Volume 1, No. 8, December 1987.
- [2] Cerf, V., "Internet Activities Board," RFC 1160, May 1990.
- [3] Vaudreuil, G., "The Federal Research Internet Coordinating Committee and the National Research Network," *ConneXions*, Volume 2, No. 10, October 1988.
- [4] Dern, D., "The ARPANET is Twenty," *ConneXions*, Volume 3, No. 10, October 1989.
- [5] Gross, P., "The Internet Engineering Task Force (IETF)," *ConneXions*, Volume 2, No. 10, October 1988.
- [6] NNSC Staff, "Profile: NSFNET," *ConneXions*, Volume 1, No. 2, June 1987.
- [7] Braun, H-W., "The new NSFNET backbone network," *ConneXions*, Volume 2, No. 12, December 1988.

MIKE ROBERTS is Vice President for Networking at EDUCOM. For the last four years, he has represented EDUCOM's six hundred members in Washington in working toward the NREN. Before joining EDUCOM, he was deputy director of the computer center at Stanford University.

A booklet which describes the HPCC initiative, entitled *Grand Challenges: High Performance Computing and Communications*, can be obtained by writing to:

Committee on Physical,
Mathematical, and Engineering
Sciences
c/o National Science Foundation
Computer and Information
Science and Engineering
1800 G Street, N.W.
Washington, DC 20550



UK Internet Protocol Connectivity Activities

by

Jon Crowcroft, University College London

Introduction

Each year, in the UK, there is a networking conference for the *Joint Academic Network* (JANET) community known as *Networkshop*. This is run by the *Joint Network Team* (JNT) which is responsible for providing network services, development and advice to the JANET community. At the 1990 *Networkshop* there was concern about a need to improve the international network services. Many international services are based on the *Internet Suite of Protocols* (IP) which are heavily used in the United States and becoming popular on the European continent. The IP protocols have been developed mainly by the United States Department of Defense. In the UK JANET community, services have been based on X.25 and the so called *Coloured Book* protocols which have provided the principal services for the last ten years. At the conference the JNT invited members of the community to advise it on the response to the requirement for IP in the provision of significantly improved networking with the rest of the international community and also on the requirement for IP within the UK JANET community itself.

DoDAG

The group tasked with providing the advice was set up in mid 1990 and called the *Department of Defense Advisory Group* (DoDAG). This reported to the JNT in November 1990 and, by the beginning of 1991, the Computer Board (which funds most networking in universities' central computing services) had accepted the JNT's recommendations, based on the DoDAG's advice, that the Internet Protocols should be provided as one of the sets of services for the community, with certain constraints.

The provision of IP represents a major change in strategy that may have consequences in all aspects of networking in the JANET community, now and in the future, and the DoDAG will continue to study these during 1991. JANET has a strategy for changing from the current set of *Coloured Book* protocols to those developed by the *International Organization for Standardization* (ISO). It is possible that the use of IP may require additions to be made to this strategy. This is of particular concern and the group will focus on the problem in collaboration with other JNT advisory and management groups.

Shoestring

Work has already commenced on a technology pilot designed to explore the operational, performance and management issues associated with the introduction of an IP service within JANET. This service is planned to be in place by October 1991. The technology pilot is called *Shoestring*. When the project was developed it was uncertain whether there would be funding in which case the project would have been undertaken with zero resources—"on a shoestring." Happily funding has been obtained, but the name lives on.

Tunnelling

Initially IP ran over dedicated lines. Recently, IP has been run over other network technologies such as X.25 [3] and DECnet [4]. The technique, known as *tunnelling*, is fairly crude and wraps IP packets in some other protocol to simulate a point-to-point link. This has allowed IP to penetrate a far larger community and allowed links to be set up at minimal cost and minimal delay.

Rather than set up a completely new network the JANET IP Service, now christened *JIPS*, will be carried over the existing JANET X.25 network by tunnelling. Although this may not give the best possible service it does remove the need for any new communications links and makes good use of the existing equipment and lines.

Definition of Service Categories

The IP network has been carefully planned to make best use of the JANET X.25 network. There will be a backbone of IP routers based on the high speed JANET connections with each institute being connected to the backbone via a local router.

In some cases IP allows the running of services for which there is no Coloured Books equivalent. The *X Window System* is an example of this. In other cases IP allows the running of services which overlap the services provided by the Coloured Books. This gives a lot of concern as there is a possibility of splitting the JANET community into two isolated groups. To prevent this, protocols have been categorised as to their level of support so that protocols where there is an adequate Coloured Book protocol and gateways to IP-based services will be discouraged and others encouraged. Therefore potential services have been categorised according to the following scheme:

- *Supported*: a recommended way of achieving a service and one which will be supported via JIPS.
- *Permitted*: neither recommended nor discouraged, (see below), but not supported.
- *Strongly discouraged*: not recommended, and also measures may be taken, either locally or nationally or both in terms of some inhibiting or preventive measures. There will be scope for experiments with such protocols.

Based on this categorisation, the following decisions have been made:

Basic services

A service of Wide Area Network (WAN) IP will be offered to institutions connected to JANET with support for The X Window System, ARPA-FTP and Telnet.

Mail

The potential use of the *Simple Mail Transfer Protocol* (SMTP) turned out to be more controversial than first expected. Some computing service mail managers are concerned over the problems and deficiencies in the quality of current mail relaying services. Others were more concerned about the possible fragmentation of the existing Coloured Books service, which is seen to be providing a high quality service. Several issues were raised:

- The quality of relaying between SMTP and Grey Book mail (used extensively in JANET).
- The effect of widespread use of mail servers, and hiding of destination systems, on the usefulness of end-to-end SMTP.
- The quality of the e-mail services between the JANET community and the Internet,
- The vulnerability of the mail community to fragmentation into non-interworking sectors.

It was concluded that these issues require further study by working groups responsible for mail. For the meantime SMTP will be strongly discouraged, and mail to and from the JANET community will continue to go through `nsfnet-relay.ac.uk`. The increasingly widespread use of the superior X.400 MTA is another reason for restricting the spread of any disjoint SMTP community.

NNTP

The *Network News Transfer Protocol* (NNTP) will be permitted for the distribution of news within JANET. However news will still enter and leave the UK via the gateway at UKC. A transition target for OSI might be provided by developments in X.400 group communication services and further study is required.

continued on next page

UK Internet Activities (*continued*)

NFS	<p>Technical difficulties in managing user and group identifiers indicate that file service and file access (e.g., using the Network File System, NFS), should be strongly discouraged until more is known about management and scalability.</p> <p>However, given the value of the facilities for obtaining archives, scope is seen for some experimental use of these services for read-only access with published, multilaterally agreed upon, user identifiers. The use of AFS or the Institutional File System may be more acceptable—this needs more study.</p> <p>Services based on <i>Remote Procedure Calls</i> (RPC) should be permitted.</p>
Transport services	<p>Use of transport level facilities (e.g., UDP and TCP) will be permitted where required. However, the use of “compound” lower layers such as the use of RFC 1006 to run OSI applications over TCP/IP will be strongly discouraged.</p>
<i>r</i> -services	<p>The Berkeley <i>r</i>-services should be actively discouraged. Institutions that use them must be made aware of the security problems. Of course, with adequate add-on authentication and privacy mechanisms (e.g., Kerberos) these services may well prove safe. However, these latter enhancements are not yet widely available.</p>
Backbone routing	<p>The question of where to place routers was discussed at some length. The conclusion was that matching the IP routing to the JANET X.25 backbone allows an incremental approach to traffic management. There are eight major JANET X.25 switch sites (NOCs) and thus eight routers will be needed. With this strategy the IP traffic can be routed onto connections that are independent of those carrying CONS traffic, if necessary at a later stage by band splitting the existing lines.</p> <p>Routers for the backbone have now been ordered. A new class B address has been obtained for this backbone network. Each institution will route to these backbone routers via its own local router.</p> <p>The backbone topology is static, with site routers forwarding external requests to a nominated backbone router. Dynamic updates are used to propagate site reachability, but not to allow for topological changes, since the backbone is (effectively) fully interconnected.</p>
International connections	<p>Connection from JIPS to the US Internet will be via the existing “Fat Pipe” connected to the JANET NOC at ULCC. Connection to other European IP networks will also be handled via the NOC at ULCC. In this case IP will be carried over the IXI pan-European X.25 network put in place as part of the COSINE project. IXI will be used to reach the appropriate gateway(s) set up as part of the RIPE initiative. [2]</p>
Network Management	<p>The <i>Simple Network Management Protocol</i> (SNMP) [1] is being promoted for diagnostic purposes, with updates installed manually via Telnet for security reasons. Management of routing information is wholly under the JANET Network Executive (they manage the X.25 network) control. They will take overall responsibility for the network with much of the day-to-day running being undertaken by the NOC at ULCC.</p> <p>For further information about JIPS contact Bob Day who is responsible for the introduction and subsequent running of the service. The address appears on the next page. There is also a mailing list associated with the work. To subscribe, send mail to: janet-ip-request@jnt.ac.uk.</p>

Joint Network Team
c/o Rutherford Appleton Laboratory
Chilton, Didcot
Oxfordshire OX11 0QX
ENGLAND
E-mail: R.Day@jnt.ac.uk

UKnet A parallel activity to the development of the JANET IP Service is the IP service offered by the *UKnet* Backbone. UKnet has been the UK backbone of the worldwide *uucp* and USENET news services for more than 10 years. It currently offers services to 670 sites of which 340 are commercial sites. In early December 1990 UKnet announced that it was planning to offer IP services in 1991. That time has now arrived and the first site was linked in early March. Three more sites were added in April with more to follow throughout 1991 and beyond.

UKnet offers two IP services, firstly over 9.6 or 64K leased lines and secondly over British Telecom "PSS Plus" closed user group X.25 service. UKnet has worked closely with JANET for many years and this co-operation will continue with IP services. This will mean that in the second half of 1991 UKnet and JANET IP sites will be able to route datagrams to sites on each others networks. For more details write to:

UKnet support Group
Computing Laboratory
University of Kent
Canterbury
Kent CT2 7NF
ENGLAND

UK Internet Consortium

The UK Internet Consortium is a new venture which aims to foster the use of TCP/IP within the UK commercial sector. For more details write to:

The UK Internet Consortium
PO Box 360
Harrow HA1 4LQ
Middlesex
ENGLAND
E-mail: ip-interest@independent.uucp

An electronic mailing list to facilitate discussion of the issues has been created and subscription requests should be sent to:

ukipnet-request@independent.uucp.

References

- [1] J. D. Case, M. S. Fedor, M. L. Schoffstall, & J. R. Davin, "Simple Network Management Protocol," RFC 1157, May 1990.
- [2] S. Goldstein & C. Michau, "Convergence of European and North American Research and Academic Networking," *ConneXions*, Volume 5, No. 4, April 1991.
- [3] D. Vair, "Components of OSI: X.25—the Network, Data Link, and Physical Layers of the OSI Reference Model," *ConneXions*, Volume 4, No. 12, December 1990.
- [4] C. Malamud, "DECnet/OSI Phase V: Real OSI or Only Selected Interfaces?," *ConneXions*, Volume 4, No. 10, October 1990.

JON CROWCROFT is a Senior Lecturer at University College London where he has been engaged in Internet related research for about 10 years. He got his BA from Cambridge some time ago, his Masters from London more recently, and intends completing his PhD before Cyberspace is fully colonized. He can be reached as jon@cs.ucl.ac.uk.

PSYCOLOQUY—An Electronic Publishing Pioneer

by Stevan Harnad, Princeton University

Introduction

Scholarly communication is currently undergoing revolutionary changes comparable to the ones that resulted from the invention of the printing press. It is now possible for scholars and scientists the world over to report and discuss new ideas and findings globally, interactively, and almost instantaneously.

Most of the world's universities and research institutions are linked together by various international electronic networks such as BITNET and the Internet (called, collectively, "The Net"). Electronic mail ("e-mail") can be sent via The Net, usually within minutes, to London, Budapest, Tel Aviv, Tokyo. But the feature that has the most remarkable potential is multiple reciprocal e-mail: Electronic groups in which every message is immediately disseminated to all members.

History

These groups first formed themselves anarchically, on various networks, the biggest of them called *USENET*, and were devoted partly to technical discussion about computers and information, the technologies that built The Net, and otherwise to "flaming": free-for-all back and forth messages by anyone, on any topic under the sun. Then groups devoted to specific topics (computers, politics, language, culture, sex) began to form, and these in turn split into "unmoderated" and "moderated" groups. Anyone with an e-mail address whose institution was connected to USENET could post to an unmoderated group and the message would automatically be sent to everyone who was "subscribed" to the group.

It was because most of the unmoderated groups were quite chaotic that the moderated groups were formed. In these, all submissions had to be channeled through a "moderator," but this was usually someone with no special qualifications or expertise, so the quality of the information on the moderated groups was still very uneven, and, with a few exceptions (principally technical discussions about computing itself), the groups were mostly havens for underinformed students and dilettantes rather than respectable scholarly forums for learned specialists in the subject matter under discussion, which by now ranged across the humanities, the social sciences and the natural sciences.

This is not far from the current status quo on The Net—a communication medium with unprecedented intellectual potential so far being used mostly as a global graffiti board for trivial pursuit in all fields other than computing itself—except that some concerted efforts are now underway to channel The Net's possibilities in a more scholarly direction. One of these projects, *PSYCOLOQUY*, is currently underway at Princeton and Rutgers Universities and its progress to date has just been accorded recognition in the annual survey of the *Library Journal* (April 15, authored by Bill Katz), which selected *PSYCOLOQUY* as one of the best new magazines of 1990.

PSYCOLOQUY

Originally initiated in 1985 by Bob Morecock of Houston University as an electronic Bulletin Board called the "BITNET Psychology Newsletter," *PSYCOLOQUY* was transformed in 1989 into a refereed electronic journal and is now sponsored on an experimental basis by the Science Directorate of the American Psychological Association.

The Co-Editor for scientific contributions is Stevan Harnad, Visiting Fellow in the Department of Psychology at Princeton University and the Co-Editor for clinical, applied and professional contributions is Perry London, Dean of the Graduate School of Applied and Professional Psychology at Rutgers University. One of *PSYCOLOQUY*'s principal scholarly objectives is to implement "peer review" on the Net in psychology and its related fields (cognitive science, neuroscience, behavioral biology, linguistics, philosophy).

Scholarly Skywriting

All contributions are refereed by a member of *PSYCOLOQUY*'s 35-member Editorial Board, but the idea is not just to implement a conventional journal in electronic form. *PSYCOLOQUY* is devoted to "Scholarly Skywriting," a radically new form of communication in which authors post to *PSYCOLOQUY* a brief account of current ideas and findings on which they wish to elicit feedback from fellow-specialists as well as experts from related disciplines the world over.

The refereeing of each original posting and each item of peer feedback on it is done very quickly, sometimes within a few hours of receipt, so as to maintain the momentum and interactiveness of this remarkable new medium, just as if each contribution were being written in the sky, for all peers to see and append to. Skywriting promises to bring the speed of scholarly communication much closer to the speed of thought, while adding to it a global scope and an interactive dimension that are without precedent in human communication, all conducted through the discipline of the written medium, monitored by peer review, and permanently archived for future reference.

The idea of "Scholarly Skywriting" is derived from a feature of a more conventional journal that Harnad has been editing for fifteen years, *Behavioral and Brain Sciences* (BBS). BBS publishes "target articles" on particularly important and controversial interdisciplinary topics together with "Open Peer Commentary" from 15–25 scholars from across specialties and around the world, followed by the author's response. Open Peer Commentary has become quite a useful and influential service in the biobehavioral sciences, but it is governed by the time constraints of conventional publication. Scholarly Skywriting in *PSYCOLOQUY* is intended for that prepublication "pilot" stage of scientific inquiry in which peer communication and feedback are still critically shaping the final outcome. Here is where the Net's speed, scope and interactiveness offer the possibility of a quantum jump for scholarly inquiry.

To subscribe

PSYCOLOQUY appears in two forms. Its USENET version, called `sci.psychology.digest`, is "gatewayed" to the Net from Princeton. Its BITNET version, formerly stored at Tulane University and archived at the University of Houston, is now at Princeton too. To subscribe (free), all you need to do is send the following one line e-mail message to `listserv@pucc.bitnet`: "sub psyc Firstname Lastname" (omitting quotes and substituting your own first and last name); the message must originate from the e-mail address at which you wish to receive *PSYCOLOQUY*. Subsequent postings are sent to `psyc@pucc.bitnet` or to `psyc@phoenix.princeton.edu`.

Back issues of *PSYCOLOQUY* are archived at Princeton and can be retrieved from any Internet e-mail address directly by a simple procedure called "anonymous FTP." Princeton also has a feature called "bitftp" that allows issues to be retrieved indirectly from BITNET by e-mail.

PSYCOLOQUY (continued)

Soon, with the help of an experimental searchable data-base called PDB, kindly lent to Princeton by Bellcore, it should be possible not only to retrieve items but to do interactive full-text searches of the *PSYCOLOQUY* archive from both BITNET and the Internet.

Conclusion

The Net is still an anarchic place. Almost all the work on *PSYCOLOQUY* so far has been donated gratis by those involved in developing it. The Co-Editors provide their services for free; Rutgers University and Princeton University provide their computing resources for free; and Bellcore has provided the data base for free. The modest subsidy from the American Psychological Association is used exclusively to pay an editorial assistant to maintain the e-mail address list and bundle the postings. All the parties involved are contributing their time and resources for one reason only: to explore and develop what they all feel is the vast potential of the Net in scholarly communication. The selection of *PSYCOLOQUY* as one of the best new magazines of 1990 is a welcome tribute to these pioneering efforts as well as to this promising new medium.

A Letter to the Editor

Dear Sir,

I read the article on the "Convergence of European and North American Research and Academic Networking" in the April issue of *ConneXions* with interest. I appreciate that the world of networking moves quickly but I would like to point out a number of inconsistencies in the article.

Firstly, the PACCOM links have changed so that both the Australian and New Zealand links no longer connect via Hawaii but rather via satellite links to FIX-WEST at NASA Ames in Mountain View, California. These links are currently 128 and 64 kbps respectively. There are plans to reconnect to Hawaii when suitable fibre is laid in 1992/3.

Last year we changed the name of the New Zealand network from *UNINET* to avoid any confusion with any of the other UNINET[T]s throughout the world. The name now used is *Kawaihiko* which is derived from the Maori (the language of the indigenous people of Aotearoa):

kawai — network of roots (of plants, trees, etc.)

hiko — lightning, electricity

Our network links the seven Universities and a number other organisations and is experiencing the same growth that other countries are experiencing in this area.

Finally, I'm sure that even an expatriate Irishman like me can speak on behalf of the many Australians and New Zealanders who might object to being relocated into Europe and/or North America without even a referendum!

*Andy Linton
Systems Programmer
Department of Computer Science
Victoria University of Wellington
New Zealand*

Information "Products" from NNSC

by Karen Roubicek, NNSC

Internet Resource Guide

The growth of NSFNET in the last few years has brought the benefits of networking to researchers at hundreds of academic, government and industrial sites. Network users have improved access to research tools, and there are greater possibilities for collaboration among members of the research community. But in order to take maximum advantage of more widespread and improved connectivity, users have to be aware of the resources that are available to them. Thus *The Internet Resource Guide* has been developed by the *NSF Network Service Center* (NNSC).

Our goal in publishing the guide is to provide a service which, by increasing the visibility of resources accessible via NSFNET and other parts of the Internet, will expose users to those facilities that will help them do their work better.

Similar resources, such as supercomputers, libraries, databases, etc., are grouped together in sections. Each resource has a separate entry that describes the resource, identifies who can use the resource, explains how to reach the local network via the Internet, and lists contacts for more information.

Electronic distribution

The guide is distributed electronically by the NNSC. Each section (new or updated) is sent in an individual message, which contains either a plain-text or a PostScript copy of a resource description. (The text in both versions is the same; the PostScript version is generally easier to read).

Those people who prefer not to receive the guide via email may ftp the chapters that are currently available via anonymous FTP from `nnsf.nsf.net` (in directory `resource-guide`). A separate mailing list is maintained to notify readers when a chapter becomes available for FTP. To be added to any of the lists, or to receive additional information about the guide, please send a message to:
`resource-guide-request@nnsf.nsf.net`.

Tour of the Internet

The NNSC, has also developed a *Tour of the Internet* in HyperCard™ format for novice network users. The stack has basic information including history, sample e-mail, FTP, and Telnet sessions, and a glossary. The Tour is intended to be a fun and easy way to learn about the Internet.

There is a "Local Info" section in the Tour. This section is a place where an organization can add information relevant to its own group of users, for example a listing of resources at that site, or other, specialized information.

System requirements

In order to run this stack, a user needs to have HyperCard 2. HyperCard 2 requires Macintosh system 6.0.5 or higher. The Tour is available via anonymous ftp on `nnsf.nsf.net`, in the directory `internet-tour`. The Internet Tour files have been compressed using *StuffIt* 1.5.1, and converted to binhex format. To use the files, the process must be reversed using the Macintosh application *StuffIt* 1.5.1. The files take up about 760k when converted to their original format. For more information about the Tour, contact the NNSC at `nnsf@nnsf.nsf.net` • 617-873-3400.

Book Review

Internetworking with TCP/IP Volume II: Design, Implementation, and Internals, by Douglas E. Comer, and David L. Stevens, Prentice-Hall, ISBN 0-13-472242-6, 1991.

The sequel

At the INTEROP show in October 1990, an announcement was made that Doug Comer would be writing a sequel to his phenomenally successful *Internetworking with TCP/IP*. Within three months, several thousand orders for the text were received, even before publication. Well, the sequel is out, selling briskly, and for good reason! Volume II of *Internetworking with TCP/IP*, picks up where Volume I left off, by examining an implementation of the protocols discussed. [Ed.: See *ConneXions*, Volume 5, No. 1, January 1991 for a review of Volume I].

Writing about an implementation is difficult, since reading code in textbook fashion is notoriously boring. Fortunately Comer and Stevens intersperse the code with lots of narrative text and numerous diagrams. As they are presenting but a single implementation strategy, they are careful in observing that there are many other possible implementation schemes, and from time to time other approaches are briefly introduced for contrast. Finally, the authors emphasize, quite rightly, that the implementation they describe is neither the least buggy, the fastest, or the most efficient—theirs is merely *one* implementation approach of several.

Xinu

The particular implementation considered is the TCP/IP code in the *Xinu* operating system, which resembles a small UNIX-like operating system (but without the procrustean licensing encumbrances of UNIX). The *Xinu* source code is openly available, and there is even a tear-sheet at end of the book with ordering information. So, the book begins with a brief introduction to the facilities provided by *Xinu* to aid in the construction of a protocol suite.

Next, the authors climb the stack, describing implementations of a device driver, ARP, IP, ICMP, UDP, and TCP, followed by the application interface to these services. Next, an implementation of a routing protocol (RIP) is described. (It should be noted that the choice of routing protocol was likely made on the basis of simplicity—RIP is a protocol whose time has come and gone, and whilst a presentation of RIP illustrates how a routing protocol interacts with IP in a system, it illustrates little else of practical value in today's internets.) Finally, the text closes with a lengthy discussion of a *Simple Network Management Protocol* (SNMP) implementation, and a brief statistical analysis of all the code discussed throughout.

From this perspective, Comer and Stevens have given us an excellent introductory text to implementing the core aspects of the Internet suite of protocols.

Problems

There are, of course, some areas which could benefit from improvement. First, as with Volume I, the text concentrates on the lower-layers of the Internet suite. It would be useful to see text on implementations of some of the applications (e.g., SMTP). Indeed, proper implementation of the applications also have profound impact on performance and usability of a system. At the very least, a discussion of the *Domain Name System* (DNS) is warranted, as this provides a basic functionality required by all systems.

Second, the section on SNMP is well-intentioned, but also somewhat misleading! The terminology used in the narrative sections is often wrong or incomplete (e.g., use of “client/server” instead of “manager/agent”), and the code presented is incorrect in many protocol functions (lucky for the authors that they didn’t write the SNMP code in Xinu). Even so, we can forgive these lapses as SNMP is one of the newer additions to the Internet suite of protocols, and hence understanding of, and experience with, SNMP is not as developed as with other members of the protocol suite.

An excellent companion

Regardless, Volume II of *Internetworking with TCP/IP* is an excellent companion to Comer’s first work, by supplementing the introduction to the Internet suite of protocols provided in Volume I.

—Marshall Rose

Call for Papers

We are looking for papers to be presented at the fifth USENIX conference on *Large Installation System Administration* (LISA), which will be held September 30 through October 3 in San Diego, California. The LISA conferences address topics of interest to people administering large UNIX sites. Previously, attempts have been made to define “large” in terms of number of machines, gigabytes of disk, or users. We feel that it is more fruitful to define a large installation as one that has problems that cannot be solved by simply scaling up well-understood solutions used on single machines with few users.

Topics

We are particularly interested in papers which take a theoretical or comparative attitude, presenting discussions of the relative merits of various approaches to problems. As always, we welcome presentation of specific solutions to specific problems, where they advance the state of the art. We also welcome papers discussing “non-technical” problems dealing with users and management. Topics of interest include but are not limited to:

- Strategies for managing data—file migration, archive systems, backup systems
- Security issues, especially where multiple people are privileged users
- Human issues of administration
- Integration of heterogeneous systems
- Usage monitoring and accounting systems
- Administration of remote sites
- Network monitoring
- Queuing systems

Papers

Papers should be from 5 to 15 pages in length, including diagrams, figures, etc. Complete papers are due by July 8, 1991; the committee will consider and comment on extended abstracts or outlines submitted by June 17th, but may require full papers to be submitted before a final decision is made. For more information, please contact the program chair:

Elizabeth D. Zwicky
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025
415-859-3290 • zwicky@erg.sri.com

Call for Presentations

NOMS '92

The IEEE 1992 *Network Operations and Management Symposium* (NOMS '92) will be held April 6–9, 1992, in Memphis, Tennessee. Subtitled "Networks Without Bounds," the symposium is sponsored by the IEEE Communications Society's Technical Committee on Network Operations and Management (CNOM), and cosponsored by the Technical Committees on Transmission Systems, Switching and Communications Software, with participation by the International Federation for Information Processing (IFIP) Working Group 6.6.

Topics

You are cordially invited to submit an original presentation (visuals with explanatory text) for consideration to be presented at the IEEE 1992 Network Operations and Management Symposium. For NOMS '92, CNOM and IFIP are working as a team to broaden the scope of the symposium. NOMS '92 will address the operations and management of public and private telecommunications and data networks, and of distributed systems. Attention will be given to trials and user experiences, operations and management of emerging network and systems technologies, and operations technologies. Areas of interest include, but are not limited to:

- Users' Perspectives and Needs
- Trials and Experiences
- Services & Operations Integration
- Models & Integrated Architectures
- Network Management Functions
- Standards & Standards Platforms
- Network Operations and Management for:
 - Distributed Systems
 - ISDN and Broadband
 - SONET/SDH Networks
 - Intelligent Networks
 - Wireless & Personal Communications

Submissions

Submissions should consist of up to ten visuals with explanatory text accompanying each one. The format should have a visual in the upper half of a page and the explanatory text in the lower half. A title page should contain the author's name, affiliation, complete address, phone, facsimile, and telex numbers, and a 200-word abstract. A formal paper for publication is not required. Seven copies of the submission in English should be sent to the Technical Program Chair:

Alan R. Johnston
 AT&T Bell Laboratories, Room 2B-071
 480 Red Hill Road
 Middletown, NJ 07748-3052
 USA
 Phone: +1 908-615-4653
 Fax: +1 908-615-5520
 Email: arj@homxb.att.com

All submissions will be reviewed for technical content and depth, quality, relevance, and originality. Accepted presentations will be published in the symposium proceedings.

Important Dates

Draft visuals and text due:	September 17, 1991
Notification of acceptance mailed:	November 15, 1991
Camera-ready visuals and text due:	February 21, 1992

Interop would like to hear from you!

Background

The INTEROP® conference and exhibition has grown from a small by-invitation-only gathering in 1986 to the industry's premier networking event. In 1990 more than 22,000 people attended INTEROP. INTEROP offers over 25 tutorials, 45 conference sessions, after-hours *Birds of a Feather* sessions (BOFs), and the only exhibition where vendors are *required* to connect to the show network and demonstrate interoperability. At INTEROP you can also see a number of special *Solutions Showcase* demonstrations in which groups of vendors cooperate to show the benefits of a particular emerging technology.

Topics

The INTEROP Technical Advisory Committee invites you to submit your suggestions, ideas and recommendations for future conference activities. We're interested in your thoughts and desires, from complete sessions to individual topics. To help us maintain our technical relevance, you can suggest tutorials, BOFs, Solution Showcase demonstrations, and the like. The more emphasis you place on *technology*, and the less of products, the more likely it is that we can incorporate your ideas. We're interested in your ideas and your participation, if appropriate. To that end, we invite you to tell us how you would like to share your expertise and experience with the INTEROP attendees. That can range from suggesting a particular theme to organizing and chairing a specific technical session, at any INTEROP event. Topics within the scope of INTEROP include, but are not limited to:

- Media issues: Ethernet, FDDI, SMDS, Frame Relay, ISDN, Token Ring, 10BaseT, Cabling alternatives, Wireless technologies
- Internetworking: Bridges, Routers, Gateways, Multi-protocol technologies, LAN-WAN Interconnection, TCP/IP, OSI
- Distributed Computing: Client-Server, Distributed databases, Remote Procedure Calls, Network file systems
- Applications: Remote login, File transfer, E-mail, Group communication, Directory systems, Graphics, The X Window System
- Security: Authentication, Authorization, Encryption, Public Key systems, Privacy, Trusted systems, Legal implications
- Network Management: SNMP, CMIP/CMIS, MIB and SMI design, Network and System administration, Protocol analysis, Network monitoring tools
- Case studies and user experiences

Suggestions

To submit your suggestion, include the following information: your name, address, phone number, e-mail, a brief outline of your topic, and your level of involvement, and send it to:

Interop, Inc.

Attn.: Executive Director, INTEROP Technical Advisory Committee

480 San Antonio Road, Suite 100

Mountain View, CA 94040-1219

USA

Phone: +1 415-941-3399 • 1-800-INTEROP (USA only)

Fax: +1 415-949-1779 • E-mail: interop@interop.com

Deadlines

We need your ideas by July 15, 1991 for our INTEROP 92 Spring event, but suggestions received after that date are also solicited and will be incorporated into future INTEROP events. (INTEROP will take place twice a year, Spring and Fall, starting in 1992). We look forward to hearing from you. With your suggestions we can make INTEROP an even more valuable experience.

CONNEXIONS

480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNEXIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President,
Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute

Subscribe to CONNEXIONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNEXIONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNEXIONS

Back issues available upon request \$15./each
Volume discounts available upon request

480 San Antonio Road, Suite 100
Mountain View, CA 94040 U.S.A.
415-941-3399 FAX: 415-949-1779

CONNEXIONS